

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 0 848 347 A1**

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
17.06.1998 Bulletin 1998/25

(51) Int. Cl.⁶: G06K 9/62

(21) Application number: 96402700.7

(22) Date of filing: 11.12.1996

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE
Designated Extension States:
AL LT LV RO SI

(72) Inventor:
The designation of the inventor has not yet been
filed

(71) Applicant: SONY CORPORATION
Tokyo (JP)

(74) Representative:
Thévenet, Jean-Bruno et al
Cabinet Beau de Loménie
158, rue de l'Université
75340 Paris Cédex 07 (FR)

(54) Method of extracting features characterising objects

(57) A method of extracting features characterising objects consists in discriminating an object from a context, the discriminating step comprising, determining the context consisting of at least two objects, choosing a topic object from the context, deriving feature sets of the topic object and the other objects in the context, judging whether there is at least one distinctive feature set in the features sets, deriving a new feature set of there is not a

distinctive feature set amongst the feature sets, and registering one of the distinctive feature sets as an outcome; and restarting the discriminating step when there exists an object which is not discriminated from the other objects. This method represents a selectionist approach to the characterisation of objects, driven by a discrimination task.




Figure 3: Feature detectors grow hierarchically as needed by the task domain.

EP 0 848 347 A1

Description

The present invention relates to the categorisation or characterisation performed by an agent (typically a robot or computer system) on the basis of perceived properties or qualities and, notably, the invention relates to the extraction of features characterising objects.

The present invention consists in a mechanism for the automatic and spontaneous formulation of "meanings", which are perceptually grounded, under selectionist pressure arising from a discrimination task.

The invention has arisen in the context of a larger research program to understand the origins of language and meaning using complex systems mechanisms such as self-organisation, co-evolution, and level formation (see "Synthesising the origins of language and meaning using co-evolution and self-organisation" by L. Steels in "Evolution of Human Language" edited by J. Hurford, Edinburgh University Press, Edinburgh, 1996c).

The invention concerns the "meaning creation" process, that is, the method whereby the agent derives characterisation rules and/or information. A method is proposed whereby an autonomous agent may originate new meanings. The agent is autonomous in the sense that its ontology is not explicitly put in by a designer, nor is there any explicit instruction.

For the purposes of the present document, meaning is defined as a conceptualisation or categorisation of reality which is relevant from the viewpoint of the agent. Meanings can be expressed through language, although they need not be.

In very general terms, meaning takes many forms depending on the context and nature of the situation concerned. Some meanings (such as colours) are perceptually grounded. Others (such as social hierarchies) are grounded in social relations. Still others (such as goals or intentions for actions) are grounded in the behavioural interaction between the agent and the environment. The present invention focuses on perceptually grounded meanings, although the proposed mechanism could also be used for other domains.

The proposed method may be employed in a wide variety of different applications where agents (e.g. software agents or robotic agents) autonomously have to make sense of their environment. This ability is especially important in the case of devices, such as unmanned exploration vehicles, which are required to navigate and perform tasks in, and/or gather data concerning, a little-known and/or hostile environment.

The description below focuses on "meaning creation" in a single agent. Work is under way to also study "meaning creation" in multiple agents and investigate how a common language can act as a way to achieve a coherent conceptual framework between agents even though every agent individually builds up its own repertoire.

There is no suggestion that the method according to the present invention corresponds empirically to any processes performed in the animal or human mind. This method does, however, enable artificial devices to create meaningful characterisation information in a wide variety of applications.

There has been a lot of work on the problem of meaning creation particularly in the connectionist literature (see, for example, "Explorations in Parallel Distributed Processing" edited by J.L. McClelland and D.E. Rumelhart, MIT Press/Bradford Books, Cambridge, Massachusetts, 1986). A perceptron, for example, can be seen as a device that acquires a set of distinctions as relevant for a classification task. The sensory channels are constituted by the inputs to the perceptron, and the weights perform the function of selecting out regions which will be input for the classification process.

The most important differences between these connectionist approaches and the present invention is that:

- (1) connectionist networks embed the build-up of a feature repertoire within the task of classification (as opposed to discrimination), and
- (2) an inductive/instructional approach as opposed to a selectionist approach is used.

An inductive approach is based on going through a (typically large) set of examples which drives the weights step-wise to reflect the best classification. In a selectionist approach, a structure comes into existence by variation or construction and is then tested as a whole for fitness in the environment. Inductive approaches result in gradual generalization. Selectionism immediately gives generalisations which might be refined more gradually.

The selectionist approach followed in the present invention is more in tune with work on feature generation in genetic algorithms research (see, for example, "Genetic Programming" by J. Koza, MIT Press, Cambridge, Ma, 1992), unsupervised learning as exemplified by the Kohonen network (see "Self-Organization and Associative Memory" by T. Kohonen, Springer Series in Information Sciences, vol.8, Springer Verlag, Berlin), and proposals, known as "Neural Darwinism", made by Edelman (see "Neural Darwinism: The Theory of Neural Group Selection" by G.M. Edelman, Basic Books, New York, 1987).

Edelman assumes that neuronal growth processes yield a primary repertoire stabilised by developmental selection, which is then subjected to experiential selection, yielding a secondary repertoire of categories. Using re-entrant maps and degeneracy, categorical perceptions of different objects can be compared and generalised to classes. Mean-

ing creation and classification are clearly distinct here. The selectionist pressure in the case considered by Edelman comes from statistical signal correlations (for the formation of classes). By way of contrast, in the present invention, the selectionist pressure comes from a discrimination task.

Below, after describing the invention in general terms, the proposed mechanisms will be explained more formally. Then some experimental results are reported. The final section of the description contains some conclusions.

Considering, first, the invention in general terms, it should be noted that agents engage in tasks relevant for their survival in a specific environment. Here the focus is on perceptually grounded discrimination tasks. The agent attempts to distinguish one object or situation from others using sensors and low-level sensory processes. The question is whether an agent is capable of developing autonomously a repertoire of features to succeed in discrimination and to adapt this repertoire when new objects are considered. A specific attempt to perform a discrimination and the subsequent adaptation of the feature repertoire is referred to below as a "discrimination game".

Let us assume that there is a set of objects, or more generally situations, which have characteristics that are sensed through sensory channels, either derived directly from sensors or from low level sensory routines. A sensory channel preferably is adapted to yield a value between 1.0 and 0.0. For example, the sensory channels could capture properties of moving objects like size, speed, average grey area, etc., or, in a more general case, internal states reflecting motivations, sensations or actuator streams. Herein, the meaning creation process is first explained abstractly without reference to specific applications. However, laboratory experiments with real mobile robots, speech, and active vision yield a sensory basis for the mechanisms proposed here.

A meaningful distinction takes the form of a "feature", which decomposes into an attribute and a value. The feature is derived by a feature detector which discretises the continuous space of one sensory channel. The feature indicates that the value of a sensory channel falls within one subregion of the space (see fig. 1). There are absolute features, such as '(colour red)', which are based on absolute values of a sensory channel for a single object, and relative meanings (such as '(speed faster)') which compare states of sensory channels for different objects. The present discussion only focuses on absolute features. A particular attribute is not necessarily relevant for each object.

The present invention makes use of the hypothesis that "meanings" can be based on construction and selection processes embedded in discrimination tasks. Each individual agent is assumed to be capable of constructing new features, i.e. new segmentations of the continuous sensory space. The process of generating diversity and variation is subjected to selection pressure coming from the discrimination task. The agent attempts to differentiate an object from a set of other objects, which constitute the context, based on the available repertoire of features and values. A discrimination may be based on one or more features grouped as a "distinctive feature set". There may be more than one possible distinctive feature set, but also none if not enough features are available. This happens either because no feature could be found to characterise the topic, or the attributes used to characterise the topic were not applicable to the other objects in the context, or a feature does not make a sufficiently fine-grained distinction. When there is no distinctive feature set, the discrimination fails and there is pressure on the agent to construct new feature detectors.

Feature detectors are refined in a hierarchical fashion and therefore form a kind of discrimination tree. The first detector divides the space up in some regions (in the examples given in the following description, always two regions). This region in its turn might then later be segmented by an additional feature detector if objects that need to be discriminated fall within the same region. Thus feature-detectors form natural hierarchies, which go as deep as required.

The set of objects among which a discrimination has to take place is assumed to be open, in the sense that new objects may enter the environment that require different or more refined features.

The present invention provides, in addition to a feature extraction method, apparatus adapted for putting this feature extraction method into practice.

Further features and advantages of the present invention will become apparent from the description below given in association with the accompanying drawings, in which:

Figure 1 is a diagram illustrating the feature perception process;

Figure 2 is a diagram illustrating the process of discrimination;

Figure 3 is a diagram illustrating how feature detectors grow as needed by the task domain;

Figure 4 is a graph indicating how success in discrimination games changes with the number of games played, in a case where the number of objects is fixed;

Figure 5 is a graph indicating how the number of features used, and the percentage of features used out of the total available, changes with the number of discrimination games played, in a case where the number of objects is fixed;

Figure 6 is a graph indicating how success in discrimination games, and the number of features used, changes with the number of games played, in a case where the number of objects increases; and

Figure 7 is a graph indicating how the total feature repertoire, and the percentage used out of the total repertoire, changes with the number of discrimination games played, in a case where the number of objects increases.

A formal description of the characterisation process according to the present invention now follows.

TERMINOLOGY

Let there be a set of objects $O = \{o_1, \dots, o_m\}$ and a set of sensory channels $S = \{\sigma_1, \dots, \sigma_n\}$, being real-valued functions over O . Each function σ_j defines a value $0.0 \leq \sigma_j(o_i) \leq 1.0$ for each object o_i .

An agent a has a set of feature detectors

$$D_a = \{d_{a,1}, \dots, d_{a,m}\}.$$

A feature detector

$$d_{a,k} = \langle p_{a,k}, V_{a,k}, \phi_{a,k}, \sigma_j \rangle$$

has an attribute name $p_{a,k}$, a set of possible values $V_{a,k}$, a function $\phi_{a,k}$, and a sensory channel σ_j . The result of applying a feature detector $d_{a,k}$ to an object o_i is a feature written as a pair $(p_{a,k}, v)$ where p is the attribute name and

$$v = \phi_{a,k}(\sigma_j(o_i)) \in V_{a,k}$$

the value.

The feature set of a for o_i is defined as

$$F_{a,o_i} = \{ \langle p_{a,k}, v \rangle \mid d_{a,k} \in D_a, d_{a,k} = \langle p_{a,k}, V_{a,k}, \phi_{a,k}, \sigma_j \rangle, v = \phi_{a,k}(\sigma_j(o_i)) \}.$$

Two features (a_1, v_1) , (a_2, v_2) are *distinctive* if

$$a_1 \neq a_2$$

and $v_1 \neq v_2$. A distinctive feature set D_{a,o_i}^C is a set of features distinguishing an object o_i from a set of other objects C .

$$D_{a,o_i}^C = \{f \mid f = (p, v) \in F_{a,o_i} \text{ and } \forall o_c \in C \text{ either } \nexists f' = (p', v') \in F_{a,o_c} \text{ with } p = p' \text{ or } \exists f' \in F_{a,o_c} \text{ with } f \text{ and } f' \text{ distinctive}\}.$$

Clearly there can be several distinctive feature sets for the same o_i and C , or none.

DISCRIMINATION GAMES

A discrimination game

$$d = \langle a, o_i, C \rangle$$

involves an agent a , a topic $o_i \in O$, and a context $C \subset O \setminus \{o_i\}$. The outcome of the game is twofold. Either a distinctive feature set could be found, $D_{a,o_i}^C \neq \emptyset$, and the game ends in success, or no such feature set could be found, $D_{a,o_i}^C = \emptyset$, and the game ends in failure.

As part of each game the repertoire of meanings is adjusted in the following way by the agent:

1.

$$D_{a,ot}^C = \phi.$$

i.e. the game is unsuccessful. This implies that there are not enough distinctions and therefore

$$\forall oc \in C, F_{a,ot} \subseteq F_{a,oc}.$$

There are two ways to remedy the situation:

(a) If there are still sensory channels for which there are no feature detectors, a new feature detector may be constructed. This option is preferred.

(b) Otherwise, an existing attribute may be refined by creating a new feature detector that further segments the region covered by one of the existing attributes.

2. $D_{a,ot}^C \neq \phi$. In a case where there is more than one possibility, feature sets are ordered based on preference criteria. The 'best' feature set is chosen and used as the outcome of the discrimination game. The record of use of the features which form part of the chosen set is augmented. The criteria are as follows:

(a) The smallest set is preferred. Thus the least number of features is used.

(b) In the case of sets of equal size, it is the set in which the features imply the smallest number of segmentations which is chosen. Thus the most abstract features are chosen.

(c) In the case of equal depth of segmentation, it is the set of which the features have been used the most which is chosen. This ensures that a minimal set of features develops.

The whole system is selectionist. Failure to discriminate creates pressure on the agent to create new feature detectors. However, the new feature detector is not guaranteed to do the job. It will be tried (next time) and only thrive in the population of feature detectors if it is indeed successful in performing discriminations.

Implementation

The discrimination game defined above has been implemented and encapsulated as an agent. The programs create a set of sensory channels and an initial set of objects which have arbitrary values for some of the sensory channels. A typical example is the following list of objects and associated values for channels:

```

o-0: [sc-3:0.73] [sc-4:0.82] [sc-5:0.07]
o-1: [sc-0:0.89] [sc-3:0.02] [sc-4:0.56] [sc-6:0.48]
o-2: [sc-0:0.74] [sc-1:0.92] [sc-2:0.22] [sc-3:0.56]
      [sc-8:0.52] [sc-9:0.03]
o-3: [sc-2:0.36] [sc-3:0.09] [sc-4:0.14]
o-4: [sc-1:0.47] [sc-2:0.61] [sc-3:0.69] [sc-5:0.67]
      [sc-6:0.14] [sc-9:0.43]
o-5: [sc-1:0.84] [sc-4:0.82] [sc-5:0.70] [sc-8:0.81]
o-6: [sc-1:0.40] [sc-2:0.32] [sc-3:0.68] [sc-4:0.96]
      [sc-5:0.41] [sc-7:0.14] [sc-8:0.76]
o-7: [sc-1:0.84] [sc-2:0.89] [sc-3:0.63] [sc-8:0.41]
o-8: [sc-0:0.72] [sc-1:0.02] [sc-3:0.92] [sc-4:0.44]
      [sc-5:0.04] [sc-7:0.29]
o-9: [sc-1:0.35] [sc-2:0.72] [sc-3:0.58] [sc-4:0.34]

```

A feature detector is a function assigning a feature-value to a certain attribute. The name of the attribute indicates its nature. It is of the form $sc_i - n_1 - \dots$ where i is the sensory channel followed by the number of segments of each consecutive segment. For example, $sc-5-2$ is the name of an attribute whose feature detector operates on $sc-5$ and divides it in 2 regions. $sc-5-2-2$ would be the name of an attribute that is a further refinement. ($sc-5-2 \vee 0$) is a feature combining this attribute with the value $v-0$.

In normal operation, the agent continuously goes through a loop performing the following activities:

1. A context is delineated. The context consists of the objects currently in the field of attention of the agent.
2. One object in this context is chosen randomly as topic.
3. The feature sets of the topic and the other objects in the context are derived.

4. An attempt is made to find possible discriminating feature sets.

We now show some typical situations for an agent a-5, which starts from no features at all. In the first game, a-5 tries to differentiate the object o-5 from o-3. The agent does not have a way yet to characterise the topic and creates a new attribute operating on sc-5.

a-5: o-5 <-> {o-3 }
 Topic: NIL
 Not enough features topic
 New attribute: sc-5-2

The next game to distinguish o-5 from o-9 and o-1 is already successful, because o-5 is again the topic. The context contains objects that do not have any response for sc-5, and thus no features can be constructed:

a-5: o-5 <-> {o-9 o-1 }
 Topic: ((sc-5-2 v-1))
 Context: (NIL NIL)
 Success: ((sc-5-2 v-1))

The next game is also successful because o-6 has value v-0 for sc-5-2, o-2 has nothing and o-5 has v-1.

a-5: o-6 <-> {o-2 o-5 }
 Topic: ((sc-5-2 v-0))
 Context: (NIL ((sc-5-2 v-1)))
 Success: ((sc-5-2 v-0))

In the following game the attributes are not sufficiently distinctive and therefore a new attribute is created. As long as there are possibilities to focus on additional sensory channels, existing attributes will not be refined. The new attribute operates on sc-3.

a-5: o-7 <-> {o-1 o-2 }
 Topic: ((sc-1-2 v-1))
 Context: (NIL ((sc-1-2 v-1)))
 No distinctive features but new one possible: (sc-2 sc-3 sc-8)
 New attribute: sc-3-2

When uncovered sensory channels are no longer available, more refined feature detectors for existing attributes start to be made. In the following example, o-0 fails to be distinguished from o-8 and o-1, even though a set of features is available to characterise each object. A refinement of the attribute operating over sc-5 is chosen.

a-5: o-0 \leftrightarrow {o-8 o-1 }
 Topic: ((sc-3-2 v-1)(sc-4-2 v-1)(sc-5-2 v-0))
 Context: (((sc-0-2 v-1)(sc-1-2 v-0)(sc-3-2 v-1)
 (sc-4-2 v-0)(sc-5-2 v-0)))
 ((sc-0-2 v-1)(sc-3-2 v-0)(sc-4-2 v-1)))
 No distinctive features but refinements possible.
 Refining attribute: sc-5-2 \Rightarrow sc-5-2-2

After a sufficient number of discrimination games the set of features stabilises. For the set of objects given above, the following is a stable discrimination tree. For each attribute the possible values are listed with their corresponding regions as well as the number of times a feature has been used.

sc-5-2:
 v-0: [0.00 0.50] 358.
 sc-5-2-2:
 v-0: [0.00 0.25] 31.
 sc-5-2-2-2:
 v-0: [0.00 0.12]
 sc-5-2-2-2-2:
 v-0: [0.00 0.06] ; v-1: [0.06 0.12] 3.
 v-1: [0.12 0.25]
 v-1: [0.25 0.50] 22.
 v-1: [0.50 1.00] 309.
 sc-1-2:
 v-0: [0.00 0.50] 651. ; v-1: [0.50 1.00] 628.
 sc-3-2:
 v-0: [0.00 0.50] 713. ; v-1: [0.50 1.00] 733.
 sc-8-2:
 v-0: [0.00 0.50] 15. ; v-1: [0.50 1.00] 8.
 sc-2-2:
 v-0: [0.00 0.50] 99. ; v-1: [0.50 1.00] 112.


```

sc-0-2:
  v-0: [0.00 0.50] ; v-1: [0.50 1.00] 42.
sc-4-2:
  v-0: [0.00 0.50] 223.
  sc-4-2-2:
    v-0: [0.00 0.25]
    v-1: [0.25 0.50] 1.
    (att a-5 sc-4 2 2 2):
      v-0: [0.25 0.37] 5.; v-1: [0.37 0.50] 5.
  v-1: [0.50 1.00] 215.
  sc-4-2-2:
    v-0: [0.50 0.75] 1.
    v-1: [0.75 1.00] 2.
  sc-4-2-2-2:
    v-0: [0.75 0.87] 5. ; v-1: [0.87 1.00] 2.
sc-6-2:
  v-0: [0.00 0.50] 2. ; v-1: [0.50 1.00]

```

We see that more abstract features, like (sc-1-2 v-0), are used more often. For some, like (sc-5-2 v-0), there is a deep further discrimination. For others, like (sc-5-2 v-1), there is none. Some features, like (sc-6-2 v-1), have not been used at all and could therefore be eliminated. Another experiment with the same objects but for a different agent a-6 yields a different discrimination tree. In one example, some sensory channels (such as sc-6) were not used, sc-4 was no longer refined, etc. Usually there are indeed many different possibilities.

When new objects enter an environment, it is preferable that the agent attempting to characterise objects in that environment should construct new distinctions if they are necessary. This is effectively what happens in preferred embodiments of the present invention. If new sensory channels become available, for example, because a new sensory routine has become active, then they will be exploited if the need arises.

Experimental Results

Fixed set of objects

Fig.4 is a graph illustrating how the discriminatory capacities of a single agent evolve as the number of discrimination games played by the agent increases, in a case where the total number of objects is fixed. In this example, there are 10 objects and 5 sensory channels.

Fig.4 shows a typical example where an agent builds up a repertoire of feature detectors starting from scratch. The graph shows the increasing discrimination success as experienced by the agent in discrimination games. It also shows the global success with the features so far, i.e. all objects are compared to all other objects only based on their features. Progress in finding more discriminatory features depends on encountering those objects that require more discrimination. Because context and topic are set probabilistically, this is not predictable.

The graph in Fig.5 shows, for the same experiment, the increasing number of features (as a percentage of the final total (22) features reached at the end of the experiment), and the percentage of features that is effectively used. We see that many features created earlier on are only gradually used and there are still many cases that have not been encountered.

Increasing the set of objects

In the next experiment (Fig.6), we start from a set of 10 objects and gradually add new objects in a probabilistic fashion, to reach a total of 50 objects. We see that the feature repertoire is extended occasionally. The average discrimination success remains close to the maximum (1.0) because new objects are only encountered occasionally and the feature detectors already constructed are general.

Fig.7 shows, for the same experiment, the relation between the total number of features that are available and the

features that are used, we see that the repertoire of features created in the beginning is used much more extensively, clearly showing

Initially, not many new features are introduced by the available repertoire is used better. Later on new features are indeed necessary.

Conclusions

The present invention provides a mechanism for the creation of perceptually grounded meaning giving a set of sensory channels and a series of objects among which discrimination has to take place. The mechanism is based on selectionist principles. There is a generator of variety and selection pressure coming from success or failure in discrimination. It has been shown that the system arrives quite rapidly at a set of possible features for discriminating objects. Most interestingly, the system remains adaptive when new objects are added or when new sensory channels become available.

The present invention provides a method for extracting features which characterise objects, and has been explained in general terms. However, it is to be understood that numerous concrete applications are possible. In typical applications, the "agent" seeking to characterise objects is a robot or a computer-based system. Also, in general, the sensory channels may be linked, for example, to sensors of physical parameters, such as to visual, auditory, or internal sensors.

The feature extraction method of the present invention provides, in effect, a means of automatically determining a set of rules which characterise a particular object and permit the object to be discriminated from the context in which the object is placed (usually, a set of objects). This characterisation can be used for a wide variety of purposes. For example, in a case where the "agent" is a robot which needs to navigate in a space which may contain other moving objects, the ability to characterise objects enables the robot to detect and track individual objects, even those of a kind not previously known to the robot, with a view, for example, to avoiding collisions.

Many applications can be envisaged where the feature extraction method of the present invention is applied as a preliminary step and then, in subsequent operation, the robot or computer system concerned sorts objects (either ones used in the original feature extraction process or new objects coming into the field of attention of the robot/system) into different groups by making use of the characterisation thus obtained. Furthermore, preferably, the characterisation is refined, as a dynamic process, as new objects enter the field of attention of the robot/computer system.

In some applications, it may be useful to send an unmanned vehicle into a hostile environment and command a computer system on-board the vehicle to put into practice the feature extraction method of the invention and to broadcast details of the best feature discriminating set that has been found. The feature discriminating set can be analysed in order to determine what different types of objects can be distinguished by the computer system on-board the vehicle. Should it be desired to remove certain types of object from the hostile environment, or to perform some special type of operation on such types of object, then an instruction identifying such types of object in terms of features which have a meaning for the computer system on-board the remote vehicle can be transmitted to that computer system.

Numerous other applications are possible, where the feature extraction method of the present invention is used automatically to identify individual members of some population and, thus, to enable gathered information to be correctly attributed to the different individual entities. This can be a powerful technique for gathering experimental data, especially in remote or hostile environments.

This feature extraction of the present invention is performed automatically and dynamically and so, as mentioned above, is particularly suited for application by devices which need to operate autonomously or semi-autonomously.

Many further concrete applications of the method according to the present invention will occur to the skilled person.

Claims

1. A method for extracting features characterizing objects comprising:

discriminating an object from a context, the step of discriminating including

- (a) determining the context consisting of at least two objects,
- (b) choosing a topic object from the context,
- (c) deriving feature sets of the topic object and the other objects in the context,
- (d) judging whether there is at least a distinctive feature set in the feature sets,
- (e) deriving a new feature set if there is not the distinctive feature set in the feature sets, and
- (f) registering one of the distinctive feature sets as an outcome; and

re-starting the step of discriminating when an object not discriminated from the other objects exists.

2. A method according to claim 1, wherein each feature of the feature set is expressed in a pair of an attribute and a value.
3. A method according to claim 1 or 2, wherein each feature of the feature set is derived by a feature detector segmenting a continuous space of a first sensory means.
4. A method according to claim 3, wherein the step of deriving a new feature set comprises:
if there is a second sensory means having no feature detector, creating a new feature detector on the second sensory means.
5. A method according to claim 3, wherein the step of deriving a new feature set comprises:
if there is no second sensory means having no feature detector, creating a new feature detector on the first sensory means.
6. A method according to any of claims 1 to 5, wherein the step of registering comprises:
judging whether there is the smallest distinctive feature set; and registering the smallest distinctive feature as the outcome if there is the smallest distinctive feature set.
7. A method according to claim 6, wherein the step of registering comprises:
judging whether there is the distinctive feature set having the smallest number of segmentation if there is not the smallest distinctive feature set; and
registering the distinctive feature set having the smallest number of segmentation as the outcome if there is the distinctive feature set having the smallest number of segmentation.
8. A method according to claim 7, wherein the step of registering comprises:
registering a feature set having been used the most as the outcome if there is not the distinctive feature set having the smallest number of segmentation.




Figure 1: Feature perception is the process of going from an object to a feature set in two steps: sensory channels contain states from sensors and sensory routines, and they are transformed into features by feature detectors.




Figure 2: Discrimination is the process of comparing two feature sets to find the discriminating features.




Figure 3: Feature detectors grow hierarchically as needed by the task domain.




Figure 4: The average success in discrimination games as well as the global success is shown on the y-axis. The number of discrimination is mapped on the x-axis (scale 1/10). All objects can be discriminated after 150 discrimination games.




Figure 5: The graph plots data for the same experiment as in fig 1. The total number of features and the percentage of features used of the total available at each time moment.




Figure 6: Graph showing a steady increase in the number of objects. The graph shows on the y-axis the number of objects (as a percentage of the total reached at the end, i.e. 50), the discriminatory success which remains close to the maximum, and the number of features (as a percentage of the total reached at the end, i.e. 35). The x-axis plots the number of discrimination games (scale 1/10).




Figure 7: Graph showing (on the y-axis) the relation between the increasing total feature repertoire and the percentage of the available repertoire that is used. The x-axis plots the number of discrimination games (scale 1/10).



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 40 2700

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	PATTERN RECOGNITION LETTERS, vol. 15, no. 11, 1 November 1994, pages 1071-1079, XP000478099 PERNUS F ET AL: "TWO-DIMENSIONAL OBJECT RECOGNITION USING MULTIREOLUTION NON-INFORMATION-PRESERVING SHAPE FEATURES" * Chapter 4 *	1	G06K9/62
A	PROCEEDINGS OF THE MIDWEST SYMPOSIUM ON CIRCUITS AND SYSTEMS, LAFAYETTE, AUG. 3 - 5, 1994, vol. VOL. 2, no. SYMP. 37, 3 August 1994, BAYOUMI M A; JENKINS W K (EDS), pages 1412-1415, XP000531910 KU K M ET AL: "ROBUST COEFFICIENT SELECTION FOR RECOGNITION OF COMPLEX CHINESE CHARACTERS" * Chapter III A *	1	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			G06K
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 13 May 1997	Examiner Sonius, M
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 (1.1.91) (P/4001)